

CLAIMS

1. A method for communication over a network, comprising:

assigning one or more doorbell addresses on a network interface adapter for use by a host processor;

writing a first descriptor to a system memory associated with the host processor, the first descriptor defining a first message to be sent over the network;

writing a command to a first one of the doorbell addresses instructing the adapter to read and execute the first descriptor;

writing a second descriptor to a second one of the doorbell addresses, the second descriptor defining a second message to be sent over the network;

responsive to the command having been written to the first one of the doorbell addresses, reading the first descriptor from the system memory using the network interface adapter, and sending the first message from the network interface adapter over the network responsive to the first descriptor; and

responsive to the second descriptor having been written to the second one of the doorbell addresses, sending the second message from the network interface adapter over the network.

2. A method according to claim 1, wherein assigning the one or more doorbell addresses comprises allocating a priority area for writing the descriptors within an address range defined by the one or more doorbell addresses, and wherein writing the second descriptor comprises writing the second descriptor to the priority area.

3. A method according to claim 2, wherein writing the second descriptor to the priority area comprises writing the second descriptor after writing the command to the first one of the doorbell addresses, and wherein sending the second message comprises, responsive to writing the second descriptor to the priority area, sending the second message before sending the first message.
4. A method according to claim 2, wherein writing the second descriptor comprises writing the second descriptor to the system memory, as well as to the priority area, and wherein sending the second message comprises, when the second descriptor is successfully written in its entirety to the priority area, executing the second descriptor written to the priority area without reading the second descriptor from the system memory.
5. A method according to claim 1, wherein writing the first and second descriptors comprises indicating first and second ranges of data to be read from the system memory for inclusion in the first and second messages, respectively, and wherein sending the first and second messages comprises reading the data from the first and second ranges responsive to the first and second descriptors.
6. A method according to claim 5, wherein reading the data comprises reading the data using direct memory access (DMA) by the network interface adapter to the system memory.
7. A method according to claim 1, wherein assigning the one or more doorbell addresses comprises assigning first and second doorbell addresses respectively to first and second processes running on the host processor, and

wherein writing the command comprises writing the command to the first doorbell address using the first process, and writing the second descriptor comprises writing the second descriptor to the second doorbell address using the second process.

8. A method according to claim 1, wherein sending the first and second messages comprises sending one or more data packets over the network for each of the messages.

9. A method according to claim 8, wherein the network comprises a switch fabric, and wherein the network interface adapter comprises a host channel adapter (HCA), and wherein writing the first and second descriptors comprises submitting work requests (WRs) for execution by the HCA.

10. A method for direct memory access (DMA), comprising:
writing a first descriptor to a system memory associated with a host processor, the first descriptor defining a first operation for execution by a DMA engine;

writing a command to a first doorbell address of the DMA engine, instructing the engine to read and execute the first descriptor;

writing a second descriptor to a second doorbell address of the DMA engine, the second descriptor defining a second operation for execution by the DMA engine;

responsive to the command written to the first doorbell address, reading the first descriptor from the system memory and executing the first descriptor using the DMA engine; and

responsive to the second descriptor having been written to the second doorbell address, executing the second descriptor using the DMA engine.

11. A method according to claim 10, wherein writing the second descriptor comprises writing the second descriptor to a priority area allocated for writing the descriptors within an address range of the doorbell addresses.
12. A method according to claim 11, wherein writing the second descriptor to the priority area comprises writing the second descriptor after writing the command to the first doorbell address, and wherein executing the second descriptor comprises, responsive to writing the second descriptor to the priority area, executing the second descriptor before executing the first descriptor.
13. A method according to claim 11, wherein writing the second descriptor comprises writing the second descriptor to the system memory, as well as to the priority area, and wherein executing the second descriptor comprises, when the second descriptor is successfully written in its entirety to the priority area, reading and executing the second descriptor written to the priority area using the DMA engine, without reading the second descriptor from the system memory.
14. A method according to claim 10, wherein writing the first and second descriptors comprises indicating first and second address ranges, respectively, in the system memory, and wherein executing the first and second descriptors comprises at least one of a scatter step, comprising conveying data from a data source to at least one of the first and second address ranges, and a gather step, comprising conveying data from at least one of the first and second address ranges to a data target.

15. A network interface adapter, for coupling a host processor to a communication network, the adapter comprising:

a range of doorbell addresses in an address space of the host processor, the range including first and second doorbell addresses;

execution circuitry, adapted to send messages over the network responsive to descriptors prepared by the host processor, the descriptors including first and second descriptors; and

a doorbell handler, which is coupled to the range of doorbell addresses so as to receive a command written by the host processor to the first doorbell address, indicating that the first descriptor has been written to a system memory associated with the host processor, the first descriptor defining a first one of the messages, and so as to receive the second descriptor written by the host processor to the second doorbell address, the second descriptor defining a second one of the messages, the doorbell handler being further coupled, responsive to the command having been written to the first doorbell address, to instruct the execution circuitry to read the first descriptor from the system memory and to execute the first descriptor so as to send the first one of the messages, and responsive to the second descriptor having been written to the second doorbell address, to pass the second descriptor to the execution circuitry and to instruct the execution circuitry to execute the second descriptor so as to send the second one of the messages.

16. An adapter according to claim 15, wherein the second doorbell address is in a priority area within the address

range, allocated for writing the descriptors thereto by the host processor.

17. An adapter according to claim 16, wherein the execution circuitry comprises a scheduler, which is adapted to determine an order of execution of the descriptors by the execution circuitry, and wherein responsive to the second descriptor having been written to the priority area, the doorbell handler is adapted to place the second descriptor in the order for execution ahead of the first descriptor.

18. An adapter according to claim 16, wherein the second descriptor is written by the host processor to the system memory, as well as to the priority area, and wherein the doorbell handler is adapted, when the second descriptor is successfully written in its entirety to the priority area, to pass the second descriptor to the execution circuitry without instructing the execution circuitry to read the second descriptor from the system memory.

19. An adapter according to claim 15, wherein the first and second descriptors indicate first and second ranges of data to be read from the system memory for inclusion in the first and second messages, respectively, and wherein the execution circuitry is adapted to read the data from the first and second ranges responsive to the first and second descriptors.

20. An adapter according to claim 19, wherein the execution circuitry comprises a gather engine, which is coupled to read the data by direct memory access (DMA) to the system memory.

21. An adapter according to claim 15, wherein the first and second doorbell addresses are assigned respectively to first and second processes running on the host processor, and wherein the command is written to the first doorbell address using the first process, and the second descriptor is written to the second doorbell address using the second process.
22. An adapter according to claim 15, wherein the execution circuitry is adapted to send the first and second messages by generating data packets to send over the network for each of the messages.
23. An adapter according to claim 22, wherein the network comprises a switch fabric, and wherein the network interface adapter comprises a host channel adapter (HCA), and wherein the first and second descriptors comprise work requests (WRs) submitted by the host processor for execution by the HCA.
24. A host channel adapter, for coupling a host processor to a switch fabric, the adapter comprising:
- a range of doorbell addresses in an address space of the host processor, the range including first and second doorbell addresses;
 - execution circuitry, adapted to generate data packets for transmission over the network responsive to work requests prepared by the host processor, the work requests including first and second work requests; and
 - a doorbell handler, which is coupled to the range of doorbell addresses so as to receive a command written by the host processor to the first doorbell address, indicating that the first work request has been written to a system memory associated with the host processor,

and so as to receive the second work request written by the host processor to the second doorbell address, the doorbell handler being further coupled, responsive to the command having been written to the first doorbell address, to pass instructions to the execution circuitry to read the first work request from the system memory and to execute a first work queue element corresponding to the first work request so as to generate the data packets called for by the first work request, and responsive to the second work request having been written to the second doorbell address, to pass a work queue element corresponding to the second work request to the execution circuitry and to instruct the execution circuitry to execute the second work queue element so as to generate the data packets called for by the second work request.

25. A direct memory access (DMA) device, comprising:
a range of doorbell addresses in an address space of a host processor, the range including first and second doorbell addresses;

a DMA engine, adapted to access a system memory associated with the host processor, responsive to descriptors prepared by the host processor, the descriptors including first and second descriptors defining respective first and second operations for execution by the DMA engine; and

a doorbell handler, which is coupled to the range of doorbell addresses so as to receive a command written by the host processor to the first doorbell address, indicating that the first descriptor has been written to the system memory, and so as to receive the second descriptor written by the host processor to the second

doorbell address, the doorbell handler being further coupled, responsive to the command having been written to the first doorbell address, to instruct the DMA engine to execute the first operation responsive to the first descriptor in the system memory, and responsive to the second descriptor having been written to the second doorbell address, to instruct the DMA engine to execute the second operation.

26. A device according to claim 25, wherein the second doorbell address is in a priority area within the address range, allocated for writing the descriptors thereto by the host processor.

27. A device according to claim 26, and comprising a scheduler, which is adapted to determine an order of execution of the operations by the DMA engine, and wherein responsive to the second descriptor having been written to the priority area, the doorbell handler is adapted to place the second operation in the order for execution ahead of the first operation.

28. A device according to claim 26, wherein the second descriptor is written by the host processor to the system memory, as well as to the priority area, and wherein the doorbell handler is adapted, when the second descriptor is successfully written in its entirety to the priority area, to pass the second descriptor to the DMA engine for execution without reading the second descriptor from the system memory.

29. A device according to claim 25, wherein the first and second descriptors indicate first and second address ranges, respectively, in the system memory, and wherein the first and second operations executed by the DMA

41400S2

engine comprise at least one of a scatter operation, comprising conveying data from a data source to at least one of the first and second address ranges, and a gather operation, comprising conveying data from at least one of the first and second address ranges to a data target.

卷之三